Multi-Agent Systems, Mechanism Design, and Institutions

Jeffrey K. MacKie-Mason

School of Information and

Dept. of Economics, U of Mich

Presented at ICE'98, Charleston S.C., October 1998

A research programme:

- Designing resource-allocating multi-agent cyber-systems
  - To **solve problems**
  - To gain **fundamental insights** on the behavior of multi-agent systems

© 1998 Jeffrey MacKie-Mason

Solving problems

- Designing interplay of **controlled and uncontrollable agents** to accomplish resource allocation goals
- E.g. (ICE'98):
  - network support for multiple QoS
  - managing computing resources
  - dynamic interaction between buyers and sellers of information goods

Harnessing multi-agent systems

- **First Welfare Theorem**: A competitive equilibrium is Pareto efficient
- **Second Welfare Theorem**: Any Pareto efficient allocation can be supported by a competitive equilibrium
- So:
  - Construct a mechanism that converges to competitive equilibrium
  - If desired, move endowments to obtain preferred allocation from Pareto set

© 1998 Jeffrey MacKie-Mason

Not too fast...

- Decentralized multi-agent systems provide **poor solutions** for distressingly fundamental problems
- Many fundamental problems have distressingly high computational **complexity**

Poor solutions

- Competitive equilibria often **don't exist**, or **aren't efficient**
  - externalities
  - transaction costs
  - increasing returns
  - integer constraints
  - asymmetric information
- **Negative results** for more general mechanisms
- E.g., for 2-person exchange, **no** mechanism exists that is always guaranteed to be
  - efficient
  - non-coercive
  - no-deficit (Myerson & Satterthwaite)

© 1998 Jeffrey MacKie-Mason

Mechanism design theory

- What is the **goal**?
  - efficiency?
  - revenue maximization?
- What are the **constraints**?
  - information
  - voluntary participation
  - assumed agent strategies are plausible
  - low computation and complexity costs
- **Solve** for set of mechanisms that achieve goal subject to constraints

© 1998 Jeffrey MacKie-Mason

Pragmatic mechanism design

- What **relaxed set of constraints** admits optimal solutions?
- Alternatively, how close do **admissible mechanisms** come to optimality?
  - How much efficiency can be achieved?
  - Or, for example, how much subsidy is needed to achieve efficiency?

Example: Application to scheduling (WWWM 1998)

- Single slot demands
  - Optima exist, and each is supported by some price mechanism
  - Ascending single good auction (ASGA):
    - not guaranteed to find optimum
    - but result deviates from optimal by bounded amount proportional to minimum bid increment (ε)
- Multi-slot problems less tractable

Pragmatic mechanisms for scheduling

Auctions for scheduling

- single good ↔ combinatorial ↔ GVA
- increasing ---------->
  - scope of concerns
  - cases in which (near) efficient outcome found
  - computational complexity

Complexity

- Multi-agent systems may be **intractably complex** for problem solving
- For many problems $comp(\varepsilon,d) = o((1/\varepsilon)^d)$
  - utility and expected utility maximization
  - finding competitive, Nash, and rational expectations equilibria
  - option valuation
    - (Rust 1997, and Traub et al. 1988)

Example: Bidding for priority

- Suppose agents bid in a "smart market" for gateway priority
  - How do they **figure out** their valuations?
    - If they care about **completion time** (T) for N-packet flow…
    - …need map from EU(T(N,t)) to reservation value for immediate service on first packet
  - **Unbounded # of paths**, with some branches of unbounded length

Re-examining the foundations

- Real economies can manage complexity…
  - What are they doing, and how?
- **Institutions** are essential
  - not all problems can be simultaneously endogenized
  - institutions provide dynamic **decomposition**
  - they allocate resources but **adapt at lower frequency**

Examples

- Why are most prices in US non-negotiable?
- Why do consumers trust credit cards over the phone (or at all!)?
- So what if it's a Federal Reserve Note?
- Other:
  - technical standards
  - conventional product configurations
  - default rules (UCC)
  - enforcement mechanisms

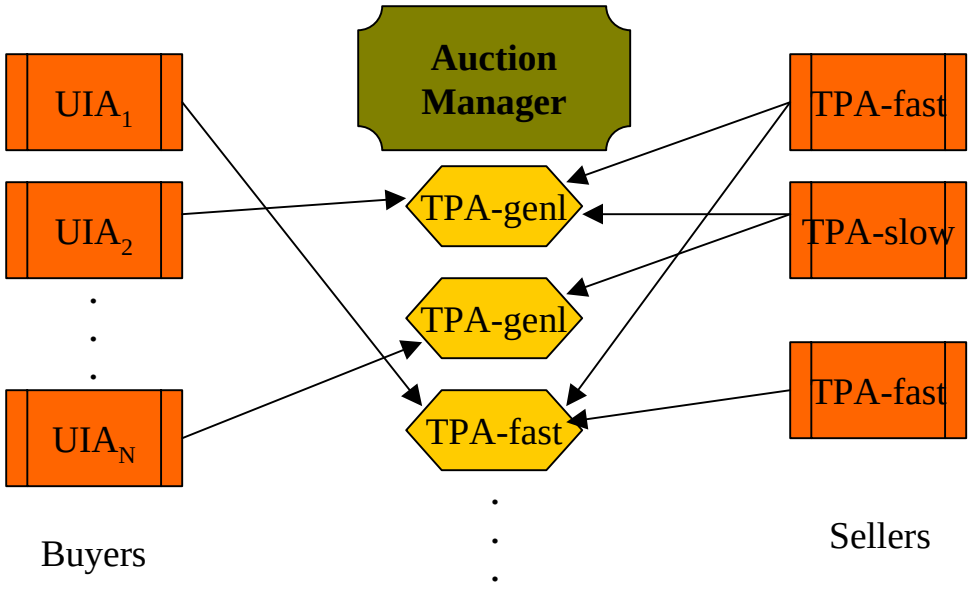
Overlapping adaptive solutions

- Long-distance phone call
  - on-demand: spot market transaction
    - spot governed by trust conventions, language, &c.
      - governed by UCC and case law
        - governed by Constitution
        - &c.
  - long-term contracts: likewise…

When is slow adaptation appropriate?

- When adaptations have **higher direct costs**
- When environments are **more stable**
- When designs are more **tightly coupled**
- When design **errors** are more consequential

MAS example: UMDL society of agents



- Who makes decisions about:
  - starting/ending auctions?
  - participating in which auctions?
  - defining scope of auctions

Which decisions are endogenous, which are not?

AMA buys disk, bandwidth from...? (Competing AMAs or central planner?)
Disk and bandwidth brokers buy resources they need from...?
Does AMA enforce allocations?
Does AMA purchase enforcement inputs from police and court agents?

© 1998 Jeffrey MacKie-Mason

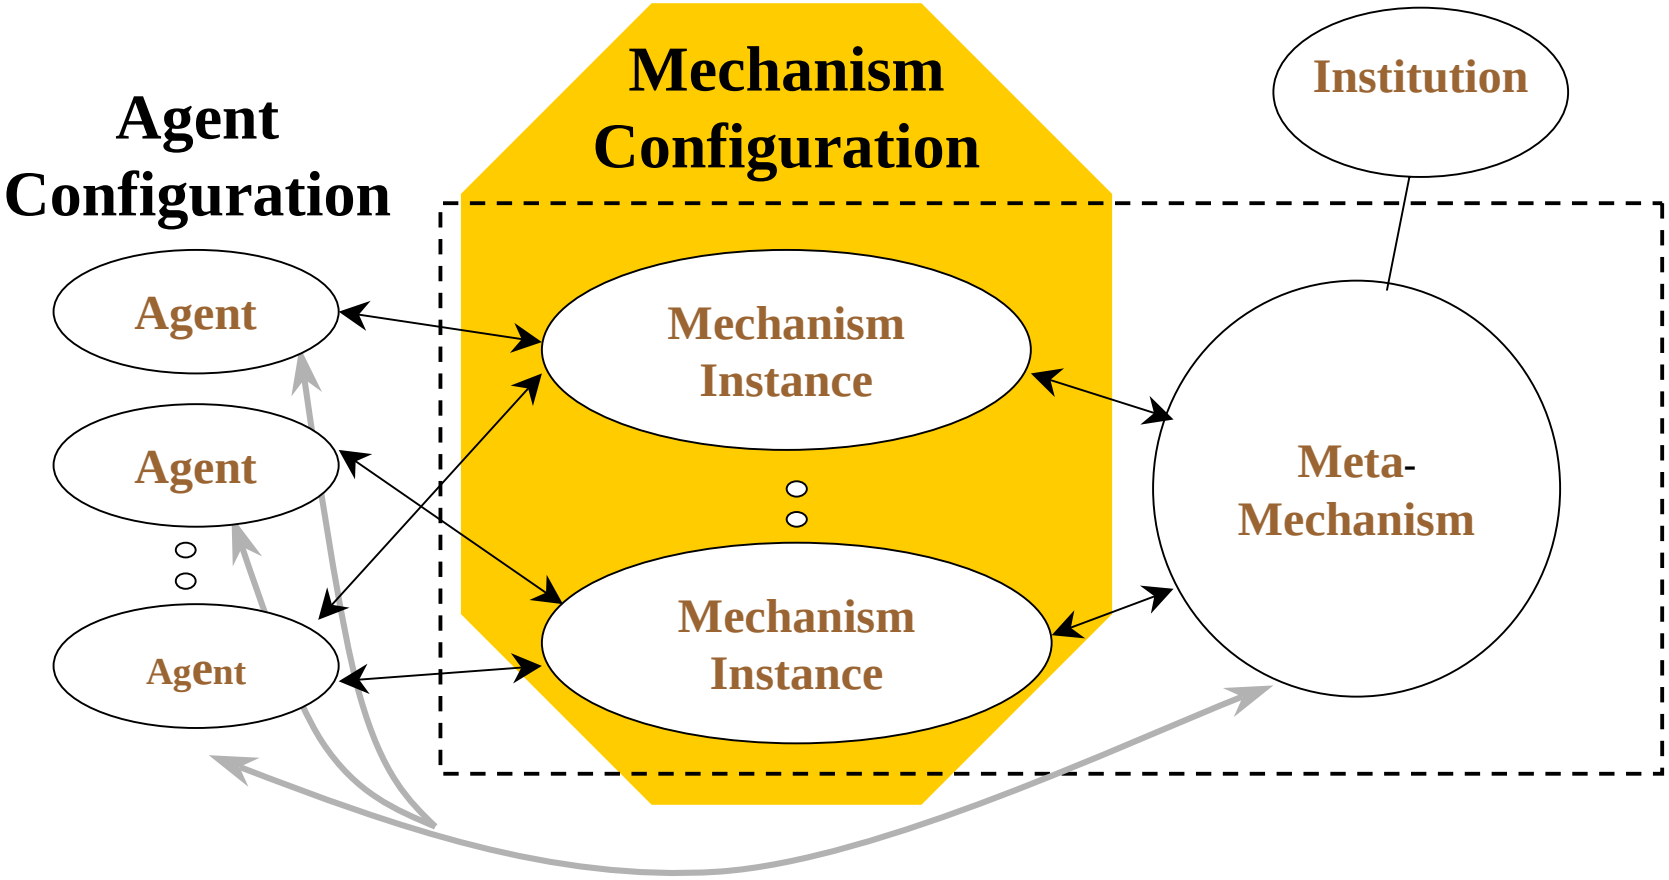
Framework: Mechanisms in multi-agent systems

- MAS need to find and settle issues that are of mutual concern to their agents
  - **which concerns** are specified?
  - **which mechanisms** are created to settle them?
- Mechanisms will generally be **approximate** and **temporary**

Endogenizing the institutions

- How do we implement the mechanism design process **internally and dynamically** within an MAS?
- Design **meta-mechanisms**: mechanisms that reconfigure mechanisms

MAS structure



Summary

- A grand adventure…
  - Cheap silicon and sand help solve hard allocation problems
  - Vast opportunities for **frictionless** commerce
- …with grand challenges
  - Existing institutions evolved for expensive computation and communications: upheaval
  - We need to understand **how to design institutions and their evolution within** multi-agent systems